Interview with Eric Allman

RIK FARROW

Eric Allman earned his BS and MS degrees from UC Berkeley in 1977 and 1980. He wrote sendmail and syslog, which became part of BSD in 1981. In 1998, Allman and Greg Olson founded Sendmail, Inc. Currently, Eric works as a Research Assistant on the Swarm project at UC Berkeley. eric.allman@gmail.com

Rik is the editor of *;login:*. rik@usenix.org

I first heard Eric Allman speak during a LISA tutorial. Eric was explaining some of the intricacies of `sendmail`, the mail server software he had written while at UC Berkeley in the early '80s.

I later cornered Eric during a conference reception, an action very unlike me. But I was determined to find out why Eric had included what I thought were three backdoors in `sendmail`, something that turned out to be incorrect. Eric also mentioned wishing he had received even a fraction of one cent for each copy of `sendmail` then in use. He later started a company that provided support for `sendmail`, a company that followed the rise and fall of the Internet boom in the late '90s.

I met with Eric in person last February in Cory Hall at the University of California, Berkeley, where he currently works. We discussed some of his past and current work.

Rik Farrow: Your experience with open source has been interesting to say the least.

Eric Allman: Open source, or if you prefer, free software, existed long before most people thought. They had IBM Share way, way back. One of the main reasons you used to go to USENIX conferences was that you always brought along six tapes with you and you walked away with six tapes, but they weren't the same six tapes you brought in. That was one of the big things about them, not just to go to talks.

RF: I believe that your open source adventure started by creating `delivermail` to handle delivery of mail that required transport beyond the local system.

EA: `delivermail` had no transport mechanisms, like `binmail`, which just delivered mail to a spool file. `delivermail` would examine the email address looking for exclamation points or at signs. If the email address didn't have these punctuations, it just appended the mail to the spool file. If the mail address did include these punctuations, then it would send the email to the correct command. Another difference between `sendmail` and `delivermail` was that `delivermail` didn't do any address translations. People had to become experts in what John Quarterman called "the matrix." One of the goals of `sendmail` was to make it easier for people to survive in this multi-network world, which included Berknet, Arpanet, and UUCP.

RF: I wanted to ask you about the backdoors in `sendmail`. When I first asked you about this many years ago, you told me you were a student maintaining `sendmail` on a small number of systems, and then someone copied `sendmail` to a machine you had no access to. The owners of that machine then demanded that you fix a bug only expressed on that system.

EA: Precisely. So I said let me log in and look at it. And they said we can't allow someone who is not part of the administrative staff onto the machine, which is normally a pragmatic approach to security. I said I will come into your office and someone can watch over my shoulder and make sure I don't do anything bad. They said, no, we can't let you on the machine. Then I can't fix your problem, and they said you have to fix our problem.

RF: A double-bind.

EA: They got more and more insistent, that I had to fix this magically somehow. And that's when the backdoor went into `sendmail`. If they won't let me on the machine, well, here's a new version, why don't we see if it fixes the problem. And it did.

The lesson out of that is the systems, including the humans that maintain them, will find a way around the security to get the job done. They actually lost security, and it would have cost them nothing to just have somebody watch me.

RF: That backdoor stayed in there for a long time.

EA: My mistake was in not taking it out immediately. The backdoor was so convenient, I thought maybe I'll leave it in and it will contribute to development. I pretty much forgot it was there.

RF: Then there was the problem with the frozen configuration file, that meant that the wizard mode password would get deleted when that was used [1].

EA: Yeah. Keep in mind that there was exactly one backdoor. There were other bugs, like ones that allowed you to clobber the stack, and you could do nefarious things there. But these were just flat out bugs.

RF: I thought that the Internet Worm used Debug, where you send a shell script as the recipient [2].

EA: Someone else put that into `sendmail`. Somebody tried to get me to put that in the `sendmail` distribution, and I said, "Are you nuts?"

RF: There's a recent movement called language security, or LangSec for short. LangSec followers believe that a key problem for most software is input parsing. It turned out that there were a lot of bugs in `sendmail` all associated with parsing, and that's because parsing is difficult.

EA: The biggest problems, of course, are buffer overflows, which are the scourge of security everywhere, and those pretty much went away after we had yet another buffer overflow and we said, "Screw it." We are just going to go around and every place we see `*p++` we are going to put a test around it.

RF: Right. In 2003, I remember that LSD had a sort of a cool exploit which wasn't a typical buffer overflow, as they figured out how write a new binary in the right place and essentially replace `sendmail` with a shell attached to an outgoing connection from port 25 [3].

EA: If I recall correctly, I had a fixed-length buffer which was pointers to opening bracket, so when I found the closing bracket I could return a pointer to the correct address.

RF: I had been single-stepping through the code, looking for where the bug was. I did find where the bug was, but by reverse engineering the patch to the source code, which of course is what hackers were doing. That's why when you said `*p++`, that brought up the memory. That was the last `sendmail` bug I remember seeing. But over the years, that whole process was very painful.

EA: Well, all I can say is `sendmail` was never as bad as Flash.

RF: Another thing you said during that short meeting, we probably only talked for 10 minutes, was if you only had a tenth of a cent for each copy of `sendmail` in use. So you eventually started `Sendmail`, Inc. Was that your idea, or did somebody approach you?

EA: I had just come out of a disastrous job, and I was sitting around, getting a little enthusiasm back, thinking what do I do next. I looked around a little bit and someone, I don't remember who, asked me if I thought about commercializing `sendmail`. I didn't know how to do that. Then I ran into a friend of mine whom I had worked for 10 years prior, and he had gone the corporate route. He helped me write a business plan and eventually agreed to come on board as their first CEO. He was a very good CEO for a company in that state. He had the sense to say at some point I'm stepping down, I like starting companies a lot more than I like running them.

RF: Those really are two different things.

EA: `Sendmail`, Inc., was a very interesting place to work, a lot going on, maybe too much. Then the Internet bubble burst [March 10, 2000]. We survived because the co-founder and I who were co-operating the company tended to be a bit more fiscally conservative than a lot of people in those days. We had a board member who said you aren't spending money fast enough, and at the very next board meeting he said you need to downsize instantly, how could you let yourself get so big. He was not my favorite board member.

RF: You survived.

EA: We did survive, but let's not go into corporate politics. `Sendmail`, Inc. got bought by Proofpoint, and the investors, including myself, got nothing out of it. But most of the employees had jobs, and the customers got taken care of. Investors, employees, customers, two out of three ain't bad. I actually was pretty happy with that.

RF: What did you do after `Sendmail`, Inc.?

EA: I kind of retired. I can afford to live as long as my tastes don't get too extravagant. That's fine, I don't have a lot of expenses. I had some offers. Then one came by email, about a new lab [4], with an invitation to come by and see what they were doing. They said they have seminars on Thursdays, including free lunch. So I

went for the free lunch. There was a research meeting right after the seminars, and I started staying for that, and at some point the person who became my supervisor said, why don't we pay you?

RF: So you are doing coding? Looks like you are involved in embedded systems here...

EA: I'm working on, loosely speaking, data storage and security for what I hesitate to call the Internet of Things, because everyone thinks they know what that means.

RF: IoT is a very broad term, from video cameras running Linux to tiny sensors with 1K of RAM...

EA: 1K? If you're lucky. There's a paper called "The Cloud Is Not Enough" [5] done by our group. These days, everyone is saying that whatever your problem is, the answer is "cloud." That's when you know you need to be looking elsewhere. The cloud may often be adequate, but there are times where it's nowhere near fast enough.

We are looking at using more distributed storage and computing. You have the cloud there, and if you have big compute jobs, you can send them off to a cloud service. If you are storing massive amounts of data, you can send them off to the cloud. We don't have objection per se to the cloud. But where somebody just unlocks a door or turns on the lights, I see no particular reason why we need to go up to the cloud and back. Our concept is that there are Swarm boxes [he indicates one sitting on the desk beside us, looking a bit like a WiFi router], and these do local processing. This box is fanless, essentially an Intel NUC; we actually have some bigger servers for storage with multiple terabytes of disk on it. Kind of a balance between the two.

RF: And there's the big problem with the Internet of Things: getting devices to play well together. And the big players have been trying to get their separate solutions accepted.

EA: Lots and lots of stovepipes.

RF: Yes, too many stovepipes. And it's *my data*, and I don't want to be sharing it to aid in marketing stuff to me.

EA: That's another point we are trying to address. It's your data, you should have access to it, you should have control over it, and we are doing security stuff: everything that goes into the db is signed, and if not marked public data, then encrypted. There are some performance issues with public key encryption, which is very slow. But that's exactly what we have implemented right now. We do a public key signature on every message that goes in, and, yes, encryption slows our system down. We have techniques for saying only every 10th record gets signed, something like that, and use hash chains. Hashes compute fast, to verify integrity. So that part's important.

References

[1] Steve Bellovin explains why the wizard mode password got deleted when a frozen configuration (sendmail.fc) file got used: http://textfiles.com/internet/sendmailwb.txt.

[2] E. H. Spafford, "The Internet Worm Program: An Analysis," Purdue Technical Report CSD-TR-823: http://spaf.cerias.purdue.edu/tech-reps/823.pdf page 5, paragraph 3.

[3] LSD sendmail reference: http://www.ouah.org/LSDsendmail.html.

[4] Swarm Lab: https://swarmlab.eecs.berkeley.edu/.

[5] B. Zhang, N. Mor, J. Kolb, D. S. Chan, N. Goyal, K. Lutz, E. Allman, J. Wawrzynek, E. Lee, and J. Kubiatowicz, "The Cloud Is Not Enough: Saving IoT from the Cloud," in *Proceedings of the 7th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud '15)*: https://www.usenix.org/system/files/conference/hotcloud15/hotcloud15-zhang.pdf.